

WHAT IS CLAIMED IS:

Sub
A1 1. A distributed file system communicating with a plurality of intelligent storage devices, wherein the distributed file system comprises a messaging system and is configured to store and manage metadata about files and directories stored on the distributed file system such that the metadata includes locations of metadata data blocks, content data blocks, and parity data blocks.

2. The distributed file system of Claim 1, further configured to provide different levels of replication and redundancy within the distributed file system, wherein the replication and redundancy information for a file or directory is stored with the file or directory's corresponding metadata.

3. The distributed file system of Claim 2, wherein the redundancy information includes parity information.

4. The distributed file system of Claim 3, wherein the parity information includes parity data blocks and the location of the parity data blocks and the metadata stores information about the location of the parity data blocks.

5. The distributed file system of Claim 1, further configured to change redundancy parameters of a file or directory, wherein the file or directory's corresponding metadata is updated to store information about the location of the redundancy data.

6. The distributed file system of Claim 5, further configured to store a copy of the file or directory data in real-time throughout a subset of the plurality of intelligent storage devices and to store the locations of the copied file or directory data in the file or directory's corresponding metadata.

7. The distributed file system of Claim 1, further configured to move data in real-time, wherein the location of the moved data is stored with the data's corresponding metadata.

8. The distributed file system of Claim 1, further configured to replicate data in real-time, wherein the location of the replicated data is stored with the data's corresponding metadata.

9. The distributed file system of Claim 8, further configured to replicate data in response to a high volume of requests for the data.

10. The distributed file system of Claim 8, further configured to replicate data in response to high utilization of the hardware which stores the data.

11. The distributed file system of Claim 8, further configure to handle more READ requests than WRITE requests.

12. The distributed file system of Claim 8, further configured to handle block transactions.

13. A virtual file system for storing data files among a plurality of modular storage units, the virtual file system comprising:

a plurality of storage units configured to:

store data blocks;

receive a data file read request;

retrieve location data information corresponding the requested data file, wherein the location data information includes storage location information about data blocks that correspond to the requested data file;

retrieve locally stored data blocks;

request remotely stored data blocks from the plurality of storage units;

receive a copy of the remotely stored data blocks from the plurality of storage units; and

return the requested data file.

14. The virtual file system of Claim 13, wherein the storage units include a storage device and a processing module, wherein the locally stored data is stored on the storage device, and the processing module retrieves the locally stored data blocks from the storage device.

15. The virtual file system of Claim 13, further comprising a write module in communication with the switch component, wherein the write module is configured to receive a data file write request and to determine the storage location of a plurality of data blocks that correspond to the data file write request.

16. The virtual file system of Claim 15, wherein the write module is further configured to distribute the plurality of data blocks among at least two of the plurality of storage units.

17. The virtual file system of Claim 16, wherein the data file write request includes mirror protection information.

5 18. The virtual file system of Claim 17, wherein the write module is further configured to distribute mirrored data that corresponds to the data file write request and complies with the mirror protection information among at least two of the plurality of storage units.

19. The virtual file system of Claim 16, wherein the data file write request includes parity protection information.

10 20. The virtual file system of Claim 19, wherein the write module is further configured to distribute parity data that corresponds to the data file write request and complies with the parity protection information among at least two of the plurality of storage units.

15 21. A modular storage unit configured to communicate with a plurality of modular storage units to provide distributed file storage, the modular storage unit comprising:

a storage module configured to store blocks of data; and

a processing module configured to:

receive requests for a data file;

20 retrieve the file location data structure corresponding to the requested data file, wherein the file location data structure includes a listing of the location of locally stored data blocks and remotely stored data blocks;

retrieve the locally stored data blocks from the storage module;

25 request remotely stored data blocks from at least one of the plurality of modular storage units;

receive the requested remotely stored data blocks; and

return the requested data file.

22. The modular storage unit of Claim 21, wherein the data blocks are of at least two different sizes.

30 23. The modular storage unit of Claim 21, wherein the data blocks are of the same size.

24. The modular storage unit of Claim 21, wherein the memory module includes a non-volatile storage device and a volatile memory device.

25. The modular storage unit of Claim 24, wherein retrieving the locally stored data blocks includes:

5 querying the volatile memory device for the requested locally stored data blocks; and

if the requested data blocks are not found, then querying the non-volatile memory device for the requested locally stored blocks.

10 26. The modular storage unit of Claim 25, wherein retrieving the locally stored data blocks includes storing the retrieved locally stored data blocks in the volatile memory device.

27. The modular storage unit of Claim 25, wherein the processing module is further configured to combine the retrieved locally stored data blocks and the received remotely stored data blocks to form the requested data file.

15 28. The modular storage unit of Claim 22, wherein the processing module is further configured to receive blocks of data for storage in the storage module.

29. A location data structure for storing information about data files and directories, wherein the data files and the directories use various redundancy and parity protection, the location data structure comprising:

20 a metadata portion which includes metadata about the corresponding data file or the data directory; and

a content portion which includes location information about the location of data file content or location data structure information about children of directories.

25 30. The location data structure of Claim 29, wherein the metadata includes at least one of name, address, type, size, security data, mirroring data, parity data, data size, block size, and location information address.

31. The location data structure of Claim 29, wherein location data structure stores information about a directory and the metadata includes mirroring data and parity data.

30 32. The location data structure of Claim 31, wherein the mirroring data is determined from the mirroring data of any descendants of the directory.

33. The location data structure of Claim 32, wherein the mirroring data matches at least the maximum mirroring of any descendants of the directory.

34. The location data structure of Claim 31, wherein the parity data is determined from the parity data of any descendants of the directory.

5 35. The location data structure of Claim 34, wherein the parity data matches at least the maximum parity protection of any descendants of the directory.

36. The location data structure of Claim 31, wherein the location information includes the location of the location data structures that correspond to the children of the directory.

10 37. The location data structure of Claim 29, wherein location data structure stores information about a data file.

38. The location data structure of Claim 37, wherein the location information includes the location of the data blocks of the file.

15 39. A smart storage device configured to communicate with other smart storage devices and to process requests for files, wherein data of the requested files is distributed among a subset of the smart storage devices.

40. The smart storage device of Claim 39, wherein the data includes of metadata and content data.

20 41. The smart storage device of Claim 39, wherein the data includes parity data.

42. The smart storage device of Claim 39, wherein the data includes mirrored data.

Add
A 7